APD: Ad Persuasion Dashboard — Insights into Facebook Political Campaign Strategies

Elyas Meguellati, Stefano Civelli, Shazia Sadiq, Gianluca Demartini
University of Queensland, Brisbane, Australia
{m.meguellati, s.civelli, s.sadiq, g.demartini}@uq.edu.au

Abstract

This demo presents an interactive dashboard for analyzing persuasive content in political social media advertising. We leverage the **APA22** dataset, comprising 56,958 unique Facebook ads from the 2022 Australian Federal Election. The system is powered by **PPAsy-DistilBERT**, a lightweight transformer model (Meguellati et al. 2026) that achieves state-of-the-art persuasion detection while requiring 86.8% less training data and up to 77.6% fewer parameters than leading ensembles. Our tool allows researchers to visualize campaign dynamics, revealing that highly persuasive ads receive 48.2% more financial investment, peak in volume by 4.8 times before election day, and utilize distinct action-oriented linguistic patterns. This system provides a scalable framework for auditing digital political campaigns.

Dashboard:
https://stefano-civelli.github.io/facebook_ads_react/
Video:
https://youtu.be/QlbkOcaEz5U

1 Introduction

Digital platforms have amplified the reach of persuasive content, necessitating robust tools to monitor political propaganda. While identifying persuasive techniques is critical for democratic transparency, the sheer scale of social media advertising renders manual analysis unfeasible. Moreover, existing computational tools often lack the granularity to detect subtle rhetorical strategies or the accessibility required by non-technical stakeholders.

To address this, we introduce an interactive dashboard powered by **PPAsy-DistilBERT**, a resource-efficient model based on DistilBERT (66M parameters). By employing a novel asymmetric binary cross-entropy loss function, our model effectively handles class imbalance across 23 persuasion techniques, outperforming larger ensembles like APatt (Purificato and Navigli 2023) (492M parameters) while using significantly less data.

The dashboard visualizes insights from the **APA22** dataset, enabling users to explore key trends in the 2022 Australian Federal Election:

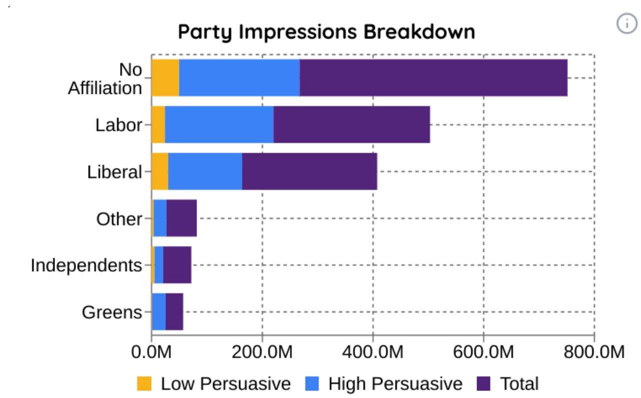

Figure 1: Party Impressions Breakdown showing distribution of lowly persuasive, highly persuasive, and total impressions for different political parties in the 2022 Australian federal election campaign.

- **Resource Allocation:** Highly persuasive ads (containing $> 80\%$ persuasive sentences) correlate with higher impact, securing 45.6% more impressions and 48.2% higher ad spend than low-persuasion content.
- **Temporal Dynamics:** Persuasive intensity escalates strategically, with the volume of high-persuasion ads peaking at 4.8 times their baseline levels immediately prior to election day.
- **Linguistic Patterns:** Lexical analysis reveals that persuasive ads prioritize broad, national themes (e.g., “future,” “better”), whereas neutral ads focus on localized, factual information.

By combining efficient NLP with intuitive visualizations, this demo offers a practical solution for unmasking persuasion strategies in modern political campaigning.

2 System Design

Data Collection. We constructed a comprehensive dataset of social media advertisements related to the Australian federal election held on May 21st, 2022. Leveraging the Meta Ad Library API[1], we systematically captured active adver-

[1]https://www.facebook.com/ads/library/

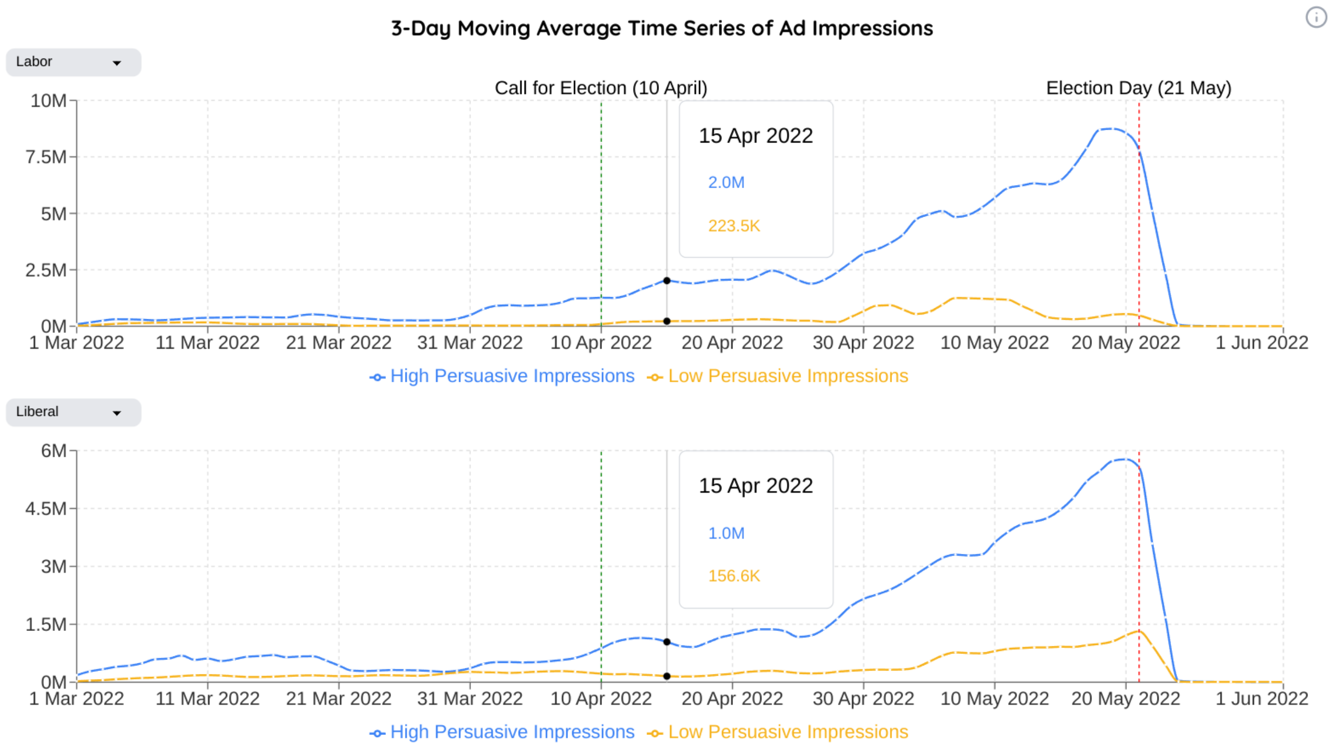

Figure 2: 3-Day Moving Average Time Series showing Impressions of Highly and Lowly Persuasive content for Labor and Liberal parties during the 2022 Australian Federal Election campaign.

tising campaigns at six-hour intervals from March 1st to June 18th, 2022. This high-frequency polling allowed us to capture short-lived advertisements that might effectively disappear from public view shortly after the election period. Each advertisement record includes detailed attributes such as creation timestamp, demographic targeting parameters (age, gender, location), ad content text, sponsoring entity, ad URL, and metrics on impressions and spending. This collection resulted in the APA22 dataset, consisting of 56,958 unique ads with an average length of 3.86 sentences per ad[2].

Data Annotation. We randomly selected 658 samples from the APA22 dataset for manual annotation of persuasive techniques. Four independent annotators labeled overlapping subsets, achieving a Fleiss' Kappa (Fleiss 1971) of 0.86. The annotated dataset was divided using stratified sampling into 75% training (493 sentences) and 25% testing (165 sentences).

Ethics Statement. The dataset collection for our study, referred to as *APA22*, underwent ethics review by authors' institution IRB. Once the data was collected, the authors were involved in the manual annotation process. The dataset we created will be made available with guidelines to prevent misuse and ensure research integrity.

Performance and Evaluation. We experimented with various transformer encoder models, including BERT, RoBERTa, and DistilBERT. While larger models offered marginal performance gains, they imposed significant latency penalties undesirable for a live web application. Our distilled model, DistilBERT (Sanh et al. 2019), offered the optimal trade-off, achieving an accuracy and F1-score of 82% with a confidence level of 89.9%. This efficiency enables our dashboard to process large batches of incoming ads with minimal delay. Furthermore, we achieved second place on the SemEval-2023 Task 3 Subtask 3 (Piskorski et al. 2023) leaderboard for persuasion detection, accomplishing this with significantly fewer computational resources compared to other top-performing teams. This validation on a standard benchmark underscores the robustness of our approach for analyzing persuasive techniques in diverse contexts.

Dashboard Architecture. Our system uses a modern web architecture designed for scalability and interactive visual analysis. The frontend is built using Next.js, a React framework that supports server-side rendering, ensuring fast load times and SEO-friendly pages. It is deployed on AWS Lambda via the Serverless Framework, allowing the application to scale automatically in response to traffic spikes during election periods.

The backend, hosted on Amazon EC2, utilizes Flask to

[2]The data collection process was reviewed by our institution's ethics committee. The dataset is available at https://docs.google.com/spreadsheets/d/1xSeorrr7qVjgnAFQ_hOU1EMPfZtqBGQQ/edit?usp=sharing&ouid=116540585721548805694&rtpof=true&sd=true.

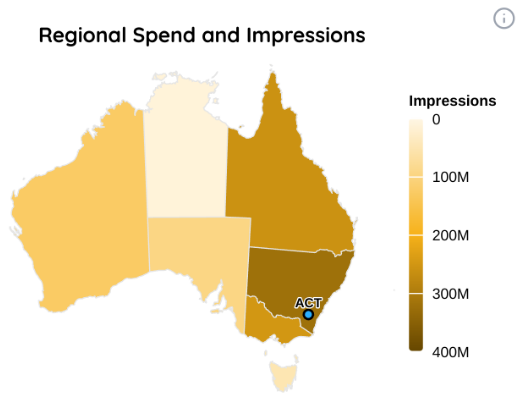

Figure 3: Regional Distribution of Ad Impressions.

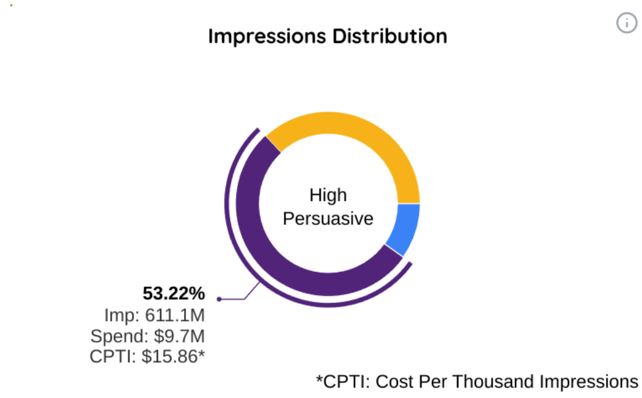

Figure 4: Impressions Distribution of High Persuasive Ads.

serve a RESTful API that ingests raw advertising data, cleans and pre-processes text, executes statistical analysis, and employs our persuasion detection model to categorize ads in real-time. Visualizations are rendered using D3.js and Recharts, providing users with interactive capabilities such as filtering by date range, political party, and persuasion intensity.

3 Persuasive Content Detection in Political Advertising

Our analysis reveals varying intensities of persuasive content in political campaigns. Our initial binary classification (0 vs 1+ persuasive sentences) proved limiting, as less than 15% of ads were categorized as non-persuasive—an expected result given that advertisements inherently aim to persuade. To address this, we refined our classification scheme based on the proportion of persuasive sentences:

- **Low Persuasion:** 20% or fewer persuasive sentences,
- **High Persuasion:** 80% or more persuasive sentences,

The 20%–80% thresholds were selected through exploratory analysis, comparing ratios such as 10%–90% and 30%–70%. The chosen split effectively distinguishes between ads with minimal and heavy reliance on persuasive techniques.

Figure 1 demonstrates this refined classification, showing party impressions in three categories. Additional insights include regional ad distribution (Figure 3) and persuasive content breakdown (Figure 4), allowing users to explore spending patterns, demographic targeting, and temporal trends for comprehensive analysis of campaign strategies.

4 Conclusion and Future Work

In this paper, we presented an interactive dashboard for analyzing persuasive content in political social media advertising. Building upon our resource-efficient **PPAsy-DistilBERT** model (Meguellati et al. 2026), we applied this tool to the **APA22** dataset, comprising over 56,000 Facebook ads from the 2022 Australian Federal Election. Our system reveals that highly persuasive ads correlate with significantly higher financial investment (+48.2% spend) and strategic temporal targeting, empowering journalists, researchers, and the public to monitor democratic processes with greater transparency.

Acknowledgments

This work is partially supported by the Australian Research Council (ARC) Centre of Information Resilience (Grant No. IC200100022) and by an ARC Future Fellowship Project (Grant No. FT240100022).

References

Fleiss, J. L. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5): 378–382.

Meguellati, E.; Civelli, S.; Bernardelle, P.; Sadiq, S.; King, I.; and Demartini, G. 2026. Towards Detecting Persuasion on Social Media: From Model Development to Insights on Persuasion Strategies. In *Proceedings of the International AAAI Conference on Web and Social Media*.

Piskorski, J.; Stefanovitch, N.; Da San Martino, G.; and Nakov, P. 2023. SemEval-2023 Task 3: Detecting the Category, the Framing, and the Persuasion Techniques in Online News in a Multilingual Setup. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, 2343–2361.

Purificato, A.; and Navigli, R. 2023. APatt at SemEval-2023 Task 3: The Sapienza NLP System for Ensemble-based Multilingual Propaganda Detection. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, 382–388.

Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

Paper Checklist

1. For most authors...
   (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
   (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
   (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
   (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? NA
   (e) Did you describe the limitations of your work? **Yes**
   (f) Did you discuss any potential negative societal impacts of your work? NA
   (g) Did you discuss any potential misuse of your work? NA
   (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes**
   (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**
2. Additionally, if your study involves hypotheses testing...
   (a) Did you clearly state the assumptions underlying all theoretical results? NA
   (b) Have you provided justifications for all theoretical results? NA
   (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? NA
   (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? NA
   (e) Did you address potential biases or limitations in your theoretical framework? NA
   (f) Have you related your theoretical results to the existing literature in social science? NA
   (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? NA
3. Additionally, if you are including theoretical proofs...
   (a) Did you state the full set of assumptions of all theoretical results? NA
   (b) Did you include complete proofs of all theoretical results? NA
4. Additionally, if you ran machine learning experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes**
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes**
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **Yes**
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **Yes**
   (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes**
   (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? NA
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
   (a) If your work uses existing assets, did you cite the creators? **Yes**
   (b) Did you mention the license of the assets? **Yes**
   (c) Did you include any new assets in the supplemental material or as a URL? **Yes**
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? NA
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? NA
   (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? **Yes**
   (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **Yes**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
   (a) Did you include the full text of instructions given to participants and screenshots? NA
   (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA
   (d) Did you discuss how data is stored, shared, and deidentified? NA